DOUBTS, INEQUALITY, AND BUBBLES

IN-KOO CHO AND KENNETH KASA

Abstract. Two agents share a common benchmark model for dividends. Each is risk-neutral but uncertainty averse, i.e., preferences are linear in consumption, but each agent has doubts about the specification of the dividend process. These doubts manifest themselves as a preference for robustness (Hansen and Sargent (2008)). Robust preferences introduce pessimistic drift distortions into the benchmark dividend process. These distortions increase with the level of wealth, and give rise to endogenous heterogeneous beliefs. Belief heterogeneity allows asset price bubbles to emerge, as in Scheinkman and Xiong (2003). A novel implication of our analysis is that bubbles are more likely to occur when wealth inequality increases. A key advantage of our analysis is that detection error probabilities can be used to assess whether empirically plausible doubts about dividends can explain observed bubble episodes.

JEL Classification Numbers: D84, G12

1. Introduction

Conventional wisdom blames the recent financial crisis on a bubble in the housing market. Bubbles have been blamed for previous crises as well. Economists do not agree on what causes bubbles. Some argue that the concept is meaningless, or that bubbles are a case of hindsight being 20/20. Others have pointed out that bubbles can be perfectly rational in a world where current outcomes depend on expectations of future outcomes. Still others argue that bubbles are *prima facie* evidence of irrationality, and suggest that economists build models in which agents suffer from various sorts of psychological biases.

Each of these approaches to bubbles has problems. Denying that bubbles exist presents the challenge of explaining why prices rise so much and fall so quickly. Although Garber (2000) points to the *possibility* that real-time assessments of fundamentals could have justified price increases in several bubble episodes, these examples do not explain their *magnitude*, nor do they explain why some people buy while others are selling, or why the bubble suddenly bursts. Theories of 'rational bubbles' show that the conditions supporting the existence of bubbles tend to be quite fragile (Santos and Woodford (1997)). Rational bubble theories also do not explain why bubbles get started in the first place, or why they are invariably correlated with large trading volumes. Finally, although models based on irrationality have the distinct advantage of being able to explain anything, the concern is that theories that explain anything explain nothing.

In our view, the most coherent and convincing account of bubbles was proposed by Scheinkman and Xiong (2003). Their model builds on the previous work of Harrison and Kreps (1978). It is based on two key ingredients: (1) Heterogeneous beliefs, and

Date: April, 2016.

We would like to thank Tom Sargent for helpful discussions.

(2) Short-selling constraints. Heterogeneous beliefs imply the coexistence of optimists and pessimists, while the presence of short-selling constraints means that the views of optimists are more fully expressed than those of pessimists. Hence, prices tend to be biased up on average, even when beliefs themselves are unbiased. What makes the model work is the fact that an agent's relative optimism fluctuates. An asset holder recognizes that at some point in the future, other traders may become relatively more optimistic. This creates an *option value* of selling to future optimists. Scheinkman and Xiong (2003) define a bubble as the value of this option, and then apply well known tools from the options pricing literature to formally calculate it as a function of the model's underlying parameters. The fact that they are able to quantify both the timing and magnitude of bubbles is what sets Scheinkman and Xiong's work apart from previous efforts to understand bubbles, because it makes their theory *testable*. For example, a key prediction of their model is that bubbles should be accompanied by large trading volumes.

Although we regard Scheinkman and Xiong (2003) as state-of-the-art when it comes to bubbles, it is not without flaws, as Scheinkman (2014) himself acknowledges in his recent Arrow lecture. One drawback concerns the underlying source of belief heterogeneity. In their model, belief heterogeneity is exogenous. Agents receive exogenous signals about future fundamentals, and are assumed to over-react to distinct signals. This makes relative optimism fluctuate. Hong and Scheinkman (2008) attempt to flesh this story out with a model of financial advisors. Still, even though by construction beliefs are unbiased on average, one could argue that the magnitude and fluctuations in relative optimism lack discipline in the sense that they are not linked to observable data (other than the bubble itself). Another drawback is that Scheinkman-Xiong abstract from learning. Over-reaction in their model is caused by agents thinking a useless signal is in fact useful. Wouldn't agents eventually learn that their signal is useless? Why don't beliefs eventually *merge* (Morris (1996))?

Our paper addresses these shortcomings. Like Scheinkman-Xiong, we assume agents are risk-neutral. However, they are *not* 'uncertainty neutral'. They have doubts about the (common) benchmark model for fundamentals. In continuous-time, these doubts are expressed as pessimistic drift distortions that are proportional to the marginal value of wealth (Hansen, Sargent, Turmuhambetova, and Williams (2006)). Hence, if wealth differs across agents, so do beliefs. Since wealth is endogenous, so are belief differences. As a corollary, an intriguing prediction of our model is that bubbles are most likely to emerge during periods of increasing wealth inequality.

Like Scheinkman-Xiong, the key mechanism in our model is belief reversals, which create an option to resell to future (relative) optimists. Although utility functions are linear, value functions are not, as they embody the option value of selling the asset in the future. In fact, the value function is convex. This implies that drift distortions *increase* with wealth. Effectively, the agent who owns the asset becomes more pessimistic as prices rise and his wealth increases. This makes sense since he has more to lose following a price decline.[1] This endogenously increasing pessimism is the source of belief reversals in our model. Although both agents are pessimistic relative to the (untrusted) benchmark model, the agent who owns the asset gets progressively more pessimistic. Eventually, he

[1]See Bhandari (2013) for a related discussion in the context of optimal risk sharing under ambiguity.

becomes even more pessimistic than the agent who sold him the asset in the first place, and he exercises his option to sell.

Besides making belief heterogeneity endogenous, another advantage of our approach is that it imparts discipline on the degree of belief heterogeneity. Following Anderson, Hansen, and Sargent (2003), we can link belief distortions to detection error probabilities. This enables us to check whether empirically plausible doubts about fundamentals can create sufficient belief heterogeneity to allow bubbles to emerge.

The remainder of the paper is organized as follows. Section 2 describes the benchmark dividend process and agent preferences. Section 3 characterizes the single-agent competitive equilibrium, using a first-order perturbation approximation of the agent's Bellman equation. With no possibility of trading, doubts about fundamentals cause asset prices to be biased below their fundamental values. Section 4 introduces a second agent. We begin by pointing out an important distinction between risk-sensitivity and robustness. In single-agent problems these two interpretations are mathematically equivalent. That is no longer the case with heterogeneous agents. If distortions are an aspect of preferences, they motivate agents to share risks. In contrast, if distortions are an aspect of beliefs, they motivate speculative trading. The two-agent competitive equilibrium now features a tension between robustness-induced pessimism, which depresses prices, and the option value of selling in the future, which raises them. In general, prices can be either above or below their Rational Expectations value. However, they are more likely to exceed their 'fundamental value' when wealth inequality increases, since this raises the resale option value. Section 5 uses simulations to explore this tension quantitatively. Using standard parameter values, we find that prices are typically below their Rational Expectations counterparts, but occasionally and persistently rise above them. This occurs when the current asset owner's relative wealth increases, which makes him relatively pessimistic. The fact that bubbles emerge when the underlying fundamentals are favorable is consistent with the evidence discussed in Garber (2000) and Scheinkman (2014). They note that bubbles typically occur in response to the introduction of new technologies or markets. In some preliminary calculations, we show that the (average) detection error probabilities associated with these bubble episodes are quite high, typically in excess of 45%. That is, the doubts that asset owners have about fundamentals would be very difficult to statistically reject, based on historical evidence. Finally, Section 6 offers a few suggestions for future work.

2. The Model

2.1. **Fundamentals.** Two risk neutral investors can trade claims on an asset. The asset is in fixed supply, normalized to one. The asset yields a stream of nonstorable dividends governed by the following geometric Brownian motion:

$$dx = \mu x \cdot dt + \sigma x \cdot dB \tag{2.1}$$

Unlike in Scheinkman-Xiong, this stream of dividends is assumed to be *observable*. Heterogeneous beliefs are instead determined by (Knightian) uncertainty about the dividend process in (2.1). That is, (2.1) is viewed as merely a useful benchmark model. Pessimistic

drift distortions reflect each investor's preference for robustness. As is typical in this literature, we assume there are no commodity or asset rental markets. Hence, if you want to consume, you have to buy the asset.

2.2. **Preferences.** Each agent has the following risk-neutral preferences:

$$V(x_0, W_0) = \max_{c,\alpha} \hat{E}_0 \int_0^\infty c e^{-rt} dt \tag{2.2}$$

where c is consumption and r is an exogeneous interest rate (assumed to be constant). The agent's budget constraint is

$$dW = [(r + \alpha(\mu_p - r))W - c]dt + \alpha W \sigma \cdot dB \tag{2.3}$$

where α is the share of wealth invested in the dividend paying asset, and μ_p is its equilibrium expected rate of return. Note that if the agent does not invest in the asset ($\alpha = 0$), his wealth just grows at the interest rate, $dW = rW \cdot dt$, since $c = 0$.

2.3. **Rational Expectations Equilibrium.** The key aspect of the preferences in eq. (2.2) is the hat over the expectations operator, $\hat{E}$. This emphasizes the fact that a preference for robustness can be thought of as producing pessimistically distorted expectations. Before getting to that, it is useful to first quickly review the nondistorted, Rational Expecations equilibrium. In this case, the dividend-paying asset is a perfect substitute for borrowing and lending, and so must offer the same rate of return. The choice of α is then a matter of indifference to the agent.

To find this rate of return, we must solve the agent's optimization problem. His Hamilton-Jacobi-Bellman (HJB) equation can be written

$$rV(x, W) = \max_{c,\alpha} \left\{ c + \mu x V_x + [(r + \alpha(\mu_p - r))W - c]V_W + \frac{1}{2}\sigma^2 x^2 V_{xx} + \frac{1}{2}\sigma^2 \alpha^2 W^2 V_{WW} + \sigma^2 \alpha x W V_{xW} \right\} \tag{2.4}$$

and the first-order condition for α is

$$\alpha = \frac{(\mu_p - r)V_W + \sigma^2 x V_{xW}}{-\sigma^2 W V_{WW}} \tag{2.5}$$

we can then sub this back into the HJB equation and solve under different assumptions about consumption. Solving PDEs is never fun, but in this case economic theory comes to the rescue. Since there is only one underlying shock, there is really only *one* state variable here. In equilibrium, the asset price and the agent's wealth will be functions of x. One can then readily verify the following solution to the HJB equation

$$V(x) = \frac{x}{r - \mu} \tag{2.6}$$

This is the Rational Expectations price of the asset, and from eq. (2.1) it implies the following equilibrium price process

$$dP = \mu P \cdot dt + \sigma P \cdot dB \tag{2.7}$$

At the same time, given risk neutrality, there are good reasons to be believe that viewed as a function of wealth, the agent's value function is just $V(W) = W$. Given the FOC in eq. (2.5) this verifies what we already knew, that in equilibrium the expected return

on the asset must be $\mu_p = r$, otherwise the agent would want to take an infinite position. That is, since the asset pays dividends, its total rate of return equals its capital gain plus its dividend yield. From the previous results:[2]

$$\mu_p = \mu + x/P = \mu + (r - \mu) = r$$

Unfortunately, we know from the work of Shiller (1989), and many others, that this is a woefully inadequate model of observed asset prices.

3. Single-Agent Equilibrium with Robustness

Let's start with the case of just one (representative) agent. Obviously there is no trading in this case. However, in contrast to the above Rational Expectations equilibrium, now suppose the agent has *doubts* about the dividend process in eq. (2.1), doubts that cannot be captured by a conventional Bayesian prior. Following Hansen and Sargent (2008), we assume the agent enlists the services of a so-called 'evil agent', who is imagined to select models so as to *minimize* the agent's expected utility. Since this worst-case model depends on the agent's own policy, the agent views himself as being immersed in a dynamic zero-sum game. To prevent the agent from being unduly pessimistic, in the sense that he would be hedging against empirically implausible models, the evil agent is assumed to pay a penalty that is proportional to the relative entropy between the evil agent's worst-cast model and the agent's own benchmark model in (2.1). It turns out that in continuous-time, these alternative models only differ in their drift terms.

To see how this works, let q_t^0 be the probability measure defined by the Brownian motion process in the benchmark model (2.1), and let q_t be some alternative probability measure, defined by some competing model. The (discounted) relative entropy between q_t and q_t^0 is then defined as follows:[3]

$$\mathcal{R}(q) = \int_0^\infty e^{-rt} \left[\int \log \left(\frac{dq_t}{dq_t^0} \right) dq_t \right] dt \tag{3.8}$$

Evidently, $\mathcal{R}(q)$ is just an expected log-likelihood ratio statistic, with expectations computed using the distorted probability measure. It can also be interpreted as the Kullback-Leibler 'distance' between q_t and q_t^0. From Girsanov's Theorem we have

$$\int \log \left(\frac{dq_t}{dq_t^0} \right) dq_t = \frac{1}{2} \hat{E} \int_0^t |h_s|^2 ds$$

where $\hat{E}$ denotes expectations with respect to the distorted measure q_t, and h_s represents a square-integrable process that is progressively measureable with respect to the filtration

[2]These results about value functions and asset prices are not quite as obvious as they seem. One might reasonably wonder what happened to the second-order partials in the above HJB equation. How did we know we could drop them? As noted by Froot and Obstfeld (1991), these higher-order terms offer the possibility of introducing nonlinear *intrinsic bubble* terms into equilibrium asset prices. Prices rise above fundamentals simply because they are expected to. However, as noted above, we do not view such rational bubbles as convincing accounts of observed bubble episodes. For one thing, they do not generate any trading volume.

[3]See Hansen, Sargent, Turmuhambetova, and Williams (2006) for a detailed discussion of robust control in continuous-time models, and in particular, on the role of discounting in the definition of relative entropy.

generated by q_t. Again from Girsanov's Theorem, we can view q_t as being induced by the following drift distorted Brownian motion[4]

$$\hat{B}(t) = B(t) - \int_0^t h_s ds$$

which then defines the following conveniently parameterized set of alternative models

$$dx = (\mu x + \sigma x h)dt + \sigma x d\hat{B} \tag{3.9}$$

and distorted budget constraint[5]

$$dW = [(r + \alpha(\mu_p - r))W - c + \alpha W \sigma h]dt + \alpha W \sigma \cdot d\hat{B} \tag{3.10}$$

The evil agent picks h subject to a (discounted) relative entropy penalty. This produces the following dynamic zero-sum game:

$$V(x_0, W_0) = \max_{c,\alpha} \min_{h} \hat{E}_0 \int_0^\infty \left(c + \frac{1}{2}\theta h^2 \right) e^{-rt} dt \tag{3.11}$$

subject to the state transition and budget equations in (3.9) and (3.10). The crucial parameter here is θ, which penalizes the actions of the evil agent. We shall see that as $\theta \to \infty$, outcomes converge to the above Rational Expectations equilibrium. The Hamilton-Jacobi-Isaacs equation for this game is

$$rV = \max_{c,\alpha} \min_{h} \left\{ c + \frac{1}{2}\theta h^2 + (\mu x + \sigma h)V_x + [(r + \alpha(\mu_p - r))W - c + \alpha W \sigma h]V_W + \frac{1}{2}\sigma^2 x^2 V_{xx} + \frac{1}{2}\sigma^2 \alpha^2 W^2 V_{WW} + \alpha \sigma^2 x W V_{xW} \right\} \tag{3.12}$$

which produces the following policy functions

$$\alpha = \frac{[(\mu_p - r) + \sigma h]V_W + \sigma^2 x V_{xW}}{-\sigma^2 W V_{WW}} \tag{3.13}$$

$$h = -\varepsilon \sigma (x V_x + \alpha W \sigma V_W) \tag{3.14}$$

where for notational convenience in what follows, we have defined $\varepsilon = \theta^{-1}$. The point to notice here is that these are game-theoretic reaction functions. The agent's optimal investment policy depends on what he thinks his evil twin is up to. At the same time, the worst-case scenario for the agent depends on his own policy function. These reaction functions allow the agent to explore in a systematic way the fragility of his policies.[6]

Although (3.13) and (3.14) are linear, and easily solved, once they are subbed back into the Bellman equation in (3.12) we obtain a nonlinear PDE that is only slightly simpler than the equations of General Relativity. Fortunately, once again, we can use economic reasoning to greatly simplify it. First, we can impose the equilibrium conditions $c = x$ and $\alpha = 1$. As in Lucas (1978), the agent's FOC for α can then be interpreted as an

[4]There are some subtleties here arising from the possibility that $\hat{B}$ and B generate different filtrations. See Hansen, Sargent, Turmuhambetova, and Williams (2006) for details.

[5]With only one underlying stochastic state variable, the evil agent is only at liberty to introduce one distortion. Hence, the h in the dividend process is the same as in the budget constraint.

[6]As in Hansen and Sargent (2008), we assume that doubts are only the mind of the agent. When simulating the model, we assume the benchmark dividend process is in fact the true data-generating process. The agent just doesn't know this. Robustness only matters because it changes the agent's policies.

equation determining the asset's equilibrium rate of return, which makes him content to hold it. Second, since again there is only one underlying state variable, we can without loss of generality drop all terms containing W, and solve for V as only a function of x.[7]

At the end of the day, we are left with the following PDE

$$rV = x + \mu x V_x + \frac{1}{2}\sigma^2 x^2 V_{xx} - \frac{1}{2}\varepsilon\sigma^2 x^2 (V_x)^2 \tag{3.15}$$

The solution determines wealth and the equilibrium price, $V(x) = W(x) = P(x)$, and Ito's Lemma can then be used to characterize their dynamics. Although much simplified, the robust HJB equation in (3.15) is still nonlinear, and as a result, does not possess a closed-form solution. Following standard practice, we therefore pursue a perturbation approximation. This is especially appealing here, since we already know the solution when $\varepsilon = 0$, so ε becomes a natural perturbation parameter.[8]

To begin, write the unknown value function as follows

$$V(x) = V^0(x) + \varepsilon V^1(x) \tag{3.16}$$

where we have dropped $O(\varepsilon^2)$ terms. We then match terms in ε. The $\varepsilon^0 = 1$ term just gives

$$rV^0 = x + \mu x V_x^0 + \frac{1}{2}\sigma^2 x^2 V_{xx}^0$$

The solution to this is easy. It's just the Rational Expectations solutions we found before[9]

$$V^0(x) = \frac{x}{r - \mu} \tag{3.17}$$

Matching the ε^1 terms gives the equation

$$\begin{aligned} rV^1(x) &= \mu x V_x^1(x) + \frac{1}{2}\sigma^2 x^2 V_{xx}^1 - \frac{1}{2}\sigma^2 x^2 (V_x^0)^2 \\ &= \mu x V_x^1(x) + \frac{1}{2}\sigma^2 x^2 V_{xx}^1 - \frac{1}{2}\sigma^2 \frac{x^2}{(r-\mu)^2} \end{aligned}$$

Note this is a *linear* ODE, with a quadratic forcing process. Solutions of the homogeneous equation generate intrinsic bubble terms, and as before we omit them.[10] Given the quadratic forcing process, it is natural to posit a particular solution of the form $V(x) = Ax^2$, where A is an undetermined coefficient. Subbing in and solving for A we find

$$V^1(x) = -\left(\frac{\sigma^2}{2(r-\mu)^2(r - 2\mu - \sigma^2)}\right) x^2$$

Combining, we get the following first-order perturbation solution to the HJB equation

$$V(x) = \frac{x}{r-\mu} - \varepsilon\left(\frac{\sigma^2}{2(r-\mu)^2(r-2\mu-\sigma^2)}\right) x^2 + O(\varepsilon^2) \tag{3.18}$$

[7]Doing this is an example of the macroeconomist's standard 'big K', 'little k' trick. Since the agent is representative, his wealth becomes a function of x.

[8]See Anderson, Hansen, and Sargent (2012) for a more sophisticated (and accurate) approach to perturbation solutions of robust control problems.

[9]Again, we are omitting nonlinear intrinsic bubble terms, which would be of the form Ax^γ here.

[10]Later, when discussing heterogeneous beliefs and speculative trading, these intrinsic bubble terms will be interpreted as resale option values, and will play a crucial role in the analysis.

This is concave as long as $r > 2\mu + \sigma^2$. In this case, notice from (3.14) that the evil agent's drift distortion decreases as dividends increase. That is, the agent's doubts dissipate as he grows wealthier. In the limit, prices converge to their Rational Expectations values. This sort of nonhomotheticity motivated Maenhout (2004) to propose a scaled version of robust control, in which $\hat{\theta} \sim \theta/V(x)$. This scaling keeps doubts alive, since the robust agent's distortion penalty effectively shrinks as the economy grows.

While homotheticity is convenient mathematically, we claim that it rules out by assumption a host of interesting and important phenomena related to the distribution of wealth.[11] For example, it seems quite reasonable that wealthier individuals would be less concerned about robustness. That's what nest eggs and rainy day funds are all about. However, if that's the case, and wealthy individuals are more willing to invest in higher yielding but more ambiguous assets, then inequality will clearly grow over time. On the other hand, one could instead argue that wealth *creates* doubts and paranoia, since wealthy people have more to lose. To quote a famous song - "Freedom's just another word for nothing left to lose".

In the following section, we consider a model where both forces are at play. The key new element is the introduction of a second agent. With nonhomothetic preferences for robustness and multiple agents, heterogeneous beliefs emerge endogenously via wealth inequality. When wealth inequality is small, doubts mainly depress asset prices, but as above, this effect will diminish as collective wealth increases. However, heterogeneous beliefs create the option to sell in the future, and like other option values, it induces *convexity* in the value function. This convexity component implies that distortions can actually increase with wealth, especially when wealth grows unequally.

4. Two-Agent Equilibrium with Robustness

The previous section considered the problem of a single agent with linear utility who wants to maximize an expected present value. The only catch was that his expectations were distorted because of doubts about the fundamentals process. In the usual way, we can then interpret the outcome as a decentralized competitive equilibrium. However, wealth dependent expectations introduces some subtleties. Notice that the agent's value function in eq. (3.18) is *concave*, despite his linear consumption preferences. This suggests that in some sense the agent is actually risk-averse, despite his linear consumption preferences. Indeed, this is the case. The HJB equation in (3.15) is the *same* HJB equation that would arise if we had endowed the agent with Duffie and Epstein's (1992) Stochastic Differential Utility preferences. In that case, the agent would have no doubts about fundamentals. He would simply dislike variance in continuation utility.

The mathematical equivalence between robustness and risk-sensitivity has been a recurring theme in the Robust Control literature. It rests on the following Legendre Transform duality

[11] For example, Bhandari (2013) studies a wide range of issues in a model with multiple agents and nonhomothetic preferences for robustness. He focuses on risk-sharing, asset pricing, and the market selection hypothesis.

$$\min_{m\geq 0, Em=1} EmV + \theta E(m\log m) = -\theta \log E \exp\left(-\frac{1}{\theta}V\right) \tag{4.19}$$

where m is a (distorted) probability measure, V is a value function, and θ is a parameter. The left-hand side provides the robust control interpretation. The hypothetical evil agent picks m subject to a relative entropy cost function. The right-hand side provides the risk-sensitivity interpretation. Here there is no evil agent, and m disappears. Optimization over m has been embedded in nonlinearity.

Although the Legendre Transform duality in eq. (4.19) points to a mathematical equivalence between robustness and risk-sensitivity, they are not *economically* equivalent. This is true even in single-agent settings. For example, Barillas, Hansen, and Sargent (2009) emphasize that the two problems call for different interpretations and calibration strategies for the θ parameter. According to Robust Control, θ is an aspect of the underlying environment. It will change if the environment changes, and its value should be calibrated using detection error probabilities. According to Risk-Sensitivity, θ is an aspect of preferences. For the usual reasons, it should then be interpreted as invariant to changes in the environment, and estimated from the data just like other preference parameters are.[12]

The economic distinction between Robust Control and Risk-Sensitivity is even more important with multiple agents, because beliefs and preferences elicit different interactions *between* agents. Belief differences motivate speculation, whereas preferance differences motivate risk-sharing. Unfortunately, even though the Robust Control literature has been active now for nearly two decades, applications with multiple agents are still hard to find. In one of the earliest examples, Anderson (2005) adopts the Risk-Sensitivity interpretation, and characterizes complete markets allocations by solving a Pareto problem. Bhandari (2013) and Borovicka (2013) adopt the Robust Control interpretation, and again use a Pareto problem to characterize complete markets allocations. A key feature of all three is that Pareto weights become endogenous state variables. In Anderson (2005) this reflects wealth dependent risk attitudes. In Bhandari (2013) and Borovicka (2013) this reflects wealth dependent beliefs. All three conclude that the distribution of wealth tends to be nonstationary with complete markets.[13]

The fact that the long-run distribution of wealth is often degenerate with complete markets and heterogeneous beliefs has sparked a recent debate about the desirability of complete markets and the suitability of the Pareto criterion when evaluating market outcomes. For example, Blume, Cogley, Easley, Sargent, and Tsyrennikov (2014) pose the following question - Suppose it's time-0 and you know that some agents will be endowed with correct beliefs, and some with incorrect beliefs. Would you vote for complete markets, or might you prefer to place restrictions on the ability to speculate? They provide examples where agents would actually vote for incomplete markets. Of course, it is crucial that this vote take place behind a veil of ignorance. Once beliefs have been endowed it is too late.

[12]Strzalecki (2011) and Maccheroni, Marinacci, and Rustichini (2006) provide more formal axiomatic linkages between robustness and recursive/risk-sensitive preferences.

[13]Borovicka (2013) is actually a bit of a hybrid. He combines (exogenous) belief distortions with recursive/risk-sensitive preferences. He finds that the distribution of wealth is more likely to be stationary with heterogeneous beliefs when preferences allow for a separation between intertemporal substitution, which governs saving, and risk aversion, which governs portfolio choice.

By definition, no one thinks their prior is wrong (although they will obviously update it in response to new information). Once priors have been endowed, agents with incorrect beliefs will continue to place immiserizing bets, forever convinced that their luck is about to change.[14]

Although certainly related, this recent literature is not directly relevant for us, for a couple of reasons. First, we do not study a complete markets general equilibrium model. As in Scheinkman and Xiong (2003), we arbitrarily shut down markets (e.g., rule out short sales), assume the existence of an exogenous risk-free asset, and do not worry about default and borrowing constraints.[15] Second, since beliefs depend on wealth in our economy, they are only truly 'heterogeneous' to the extent that initial wealth allocations are heterogeneous.

What is important for us is the economic distinction between risk aversion and belief distortions. If the agents in our economy are viewed as being risk averse, they will have an incentive to pool risk. Diversification will generally produce interior solutions to their portfolio problems. In contrast, heterogeneous beliefs (combined with linear consumption preferences) will generally produce corner solutions, with the most optimistic agent absorbing all available shares of the ambiguous asset, and its market clearing price being determined by his beliefs.[16]

NEEDS MORE DISCUSSION

With two agents and heterogeneous beliefs, the equivalence between an agent's wealth and the value of the asset breaks down. Individual wealth trajectories depend on portfolio decisions, which depend on beliefs (which, in turn, depend on wealth). As a result, our earlier strategy of simplifying the agent's HJB equation by imposing market-clearing and dropping one of the state variables will apparently no longer work. However, given that agents choose to invest all their wealth in either the risk-free asset or the higher yielding risky asset, we can still exploit our earlier results using the following two-step strategy. First, we can calculate the buy-and-hold price given an agent's current level of wealth. Then, we can simply add on the option value of selling to the other agent. This second component represents the solution to the homogeneous part of the agent's HJB equation.

Consider an agent with wealth W who chooses to hold the asset. Imposing market-clearing ($\alpha = 1$) in the FOC's in eqs. (3.13)-(3.14) and then substituting back into the HJB eq. (3.12) delivers the following PDE

$$rV = \mu x V_x + \frac{1}{2}\sigma^2 x^2 V_{xx} + rWV_W - \frac{1}{2}\sigma^2 W^2 V_{WW} + \frac{1}{2}\varepsilon\sigma^2(V_W^2 - V_x^2) \tag{4.20}$$

[14]Brunnermeier, Simsek, and Xiong (2014) and Gilboa, Samuelson, and Schmeidler (2014) propose closely related modifications of the Pareto criterion in settings with heterogeneous beliefs.

[15]Actually, given the restriction on short sales, the inattention to bankruptcy and borrowing constraints is not really an issue, given that our economy grows over time.

[16]Although we are using singular pronouns when discussing the two agents, it is important to keep in mind that we are actually considering two agent *types*, each consisting of a large number of identical agents, so that we can continue to rely on competitive market forces to determine prices and allocations.

As before, this equation is nonlinear, and does not possess a closed-form solution. However, notice that when $\varepsilon = 0$ we obtain the obvious solution $V(W) = W$. So once again we can pursue a perturbation approximation. Following the same steps as before delivers[17]

$$\begin{aligned} V(W) &= W - \frac{\varepsilon}{2}\left(\frac{\sigma^2}{r-\sigma^2}\right)W^2 + \varepsilon B W^{\beta} \\ &= \frac{x}{r-\mu} - \frac{\varepsilon}{2}\left(\frac{\sigma^2}{r-\sigma^2}\right)W^2 + \varepsilon B W^{\beta} \end{aligned}$$

where the second equation follows from the fact that $V = x/(r-\mu)$ when $\varepsilon = 0$. The equilibrium price and owner's wealth are then determined by the fixed point condition $P(x) = W(x) = V(W(x))$. The only difference here arises from the last term, which represents the solution of the homogeneous part of the $V^1(W)$ ODE. It captures the option value of selling the asset.[18] The β exponent is a root of the characteristic equation, $\frac{1}{2}\sigma^2\beta(\beta-1) - r\beta + r = 0$. The two roots are $\beta_{1,2} = (1, 2r/\sigma^2)$. In what follows, we assume $2r > \sigma^2$. In this case, the last term imparts a *convexity* component into the agent's value function.

Equation (4.21) nicely summarizes the opposing forces that arise when agents confront ambiguity, but at the same time realize that others confront ambiguity as well, which creates valuable trading opportunities. The middle term represents the effects of ambiguity for a Robinson Crusoe, who must face his doubts alone. This term depresses asset values. The last term represents the option value of selling to someone who becomes less doubtful than you. This increases asset values. Following Scheinkman and Xiong (2003), we interpret this as a *bubble*. In contrast to their work, however, it is not at all clear whether the combined forces produce asset prices that exceed their Rational Expectations values.

As usual in option pricing problems, the constant of integration, B, and the option exercise threshold are determined by the following pair of value-matching and smooth-pasting conditions

$$V^o = V^{no} - \tau \tag{4.21}$$

$$V^o_W = V^{no}_W \tag{4.22}$$

where τ represents a fixed transactions cost, and (V^o, V^{no}) denote the value functions of the current asset owner and nonowner, respectively. These conditions ensure that when one agent wants to sell, the other wants to buy, and in addition, that there is no benefit to waiting just a little bit longer to see what happens. The only tricky part is that the exercise threshold is not a fixed barrier, but rather a *manifold* in (W_o, W_{no}) space. This simply reflects the fact that option values depend on belief differences, which depend on wealth differences. A given difference of opinion can arise for a multitude of different wealth levels. To resolve this indeterminancy, we can view the owner's wealth level that triggers a sale as a function of the nonowner's wealth level. This makes the integration

[17] As before, under a buy-and-hold strategy there is really only one state variable. Hence, we can drop the terms involving x.

[18] We do not need to consider the homogeneous solution of the $V^0(W)$ ODE, because when $\varepsilon = 0$, there are no belief differences, and so no trading or option values arise.

constant, B, a function of W_{no} as well. Following this approach, the smooth-pasting condition delivers the following expression

$$B = \frac{2\phi}{\beta}\bar{W}_o^{2-\beta}$$

where for notational convenience we have defined $\phi = \sigma^2/(r-\sigma^2)$. Notice that B depends on W_{no} because the owner's selling threshold, $\bar{W}_o$, depends on W_{no}. Substituting this into the value-matching condition then delivers the following polynominal for $\bar{W}_o(W_{no})$

$$\left(1-\frac{2}{\beta}\right)\bar{W}_o^{\beta} - \left(W_{no}^2 + \frac{\tau}{\varepsilon\phi}\right)\bar{W}_o^{\beta-2} + \frac{2}{\beta}W_{no}^{\beta} = 0$$

Figure 1 depicts the value functions and optimal selling threshold when $r = .06$ and $\sigma^2 = .04$. In this case, $\beta = 3$ and the equation determining the threshold becomes a cubic. The nonowner's wealth has been held fixed at $W_{no} = 20$. Notice that his valuation is less than this due to doubts about the dividend process ($\varepsilon = .01$).

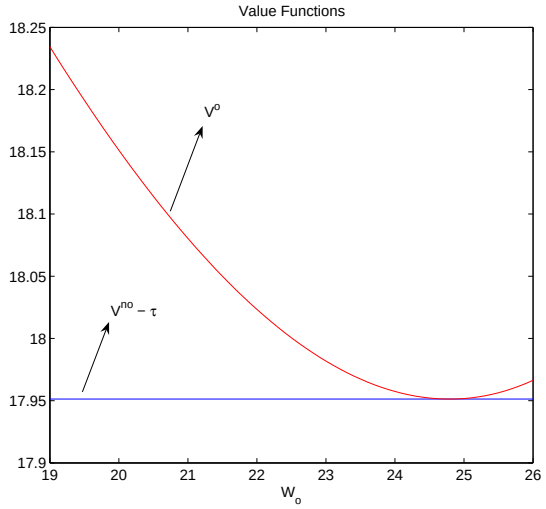

FIGURE 1. Selling Threshold: $W_{no} = 20$, $\varepsilon = .01$, $\sigma^2 = .04$, $r = .06$, $\tau = .2$

The flat line plots the nonowner's valuation net of the transactions cost. When $W_o < W_{no}$, the current owner values the asset more. However, once W_o begins to rise above W_{no}, the owner's relative pessimism *increases*, since he becomes increasingly worried about the resale value of his asset. Eventually, when W_o approaches 25, his relative pessimism is sufficient to offset the transactions cost, and a sale occurs.

5. Simulations

As noted in the Introduction, an important advantage of our robust resale option value theory of bubbles is that it makes *quantitative* predictions about the magnitude and timing of bubble episodes. This means that in principle it can be tested against the data. Whether this story is convincing will depend on whether the agents within the model are displaying empirically plausible doubts about the fundamentals process. This can be assessed by computing detection error probabilities (Anderson, Hansen, and Sargent (2003)). If these are small, then the worst case scenario that motivates the agent's behavior could be statistically rejected by the historical data. This would make the results unconvincing, as they would rely on imparting an undue degree of pessimism to the agents.

Unfortunately, we are still not quite there yet, so for now we content ourselves with performing simulations. Figures 2-4 report the results of a few representative simulations. The parameter values are the same in all three, and the implicit time unit is a year. The drift in the fundamental is 1% ($\mu = .01$) and the variance is 3% ($\sigma^2 = .03$). The risk-free rate is set to $r = .045$, and the transaction costs is set to $\tau = 0.2$. Note that the severity of this cost will depend units. We initialize fundamentals at $x_0 = 3$, so that the initial Rational Expectations value is $P = 200$. This implies a very small transaction cost, on the order of 0.1%.

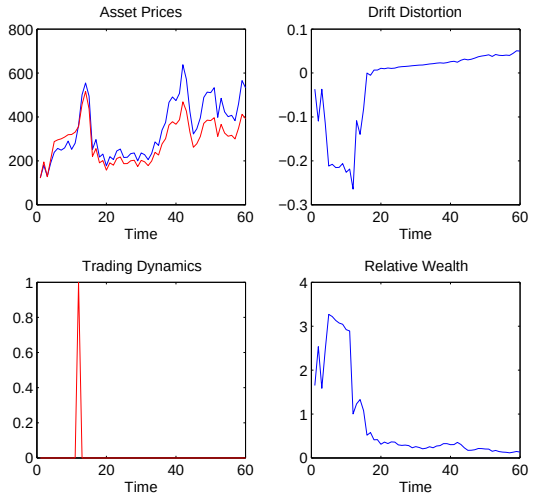

Figure 2. Simulation: $\mu = .01$, $\sigma^2 = .03$, $\varepsilon = .001$

Figure 2 displays two key features of model. First, robust values can be either above or below their Rational Expectations values. The blue line plots the Rational Expectations value, and the red line plots the robust valuation. There are two ways to define a bubble here. The first is in reference to the Rational Expectations value. By this definition, the 'bubble' in the early years of Figure 2 is quite modest. This is because doubts are depressing asset values below their Rational Expectations value. The other way is to follow Scheinkman and Xiong (2003), and define a bubble as the value of the resale option. With

this definition, bubbles would be more significant. The second key feature revealed in Figure 2 is the interplay between inequality and bubbles. The bottom right panel displays the ratio of the owner's wealth to the nonowner's wealth. Notice that around $T = 15$, the owner's relative wealth rises in response to favorable fundamentals. This starts to make him worried about the resale value of his asset. He becomes so worried that the evil agent's drift distortion exceeds 20%! In response, he decides to sell the asset.

Is this a reasonable outcome? Following Anderson, Hansen, and Sargent (2003), the Chernoff rate function is $\rho(x) = \frac{1}{8}|h(x)|^2$. We can then compute a *very* rough detection error probability using the bound

$$\text{avg det error prob} \leq \frac{1}{2}e^{-\rho T}$$

where ρ is the mean-squared drift distortion. This turns out to be .47 in Figure 2, basically because the *average* drift distortion is quite small. We need to do more work to calculate a tighter, state-dependent, bound. This would allow us to assess whether doubts during the bubble are reasonable.

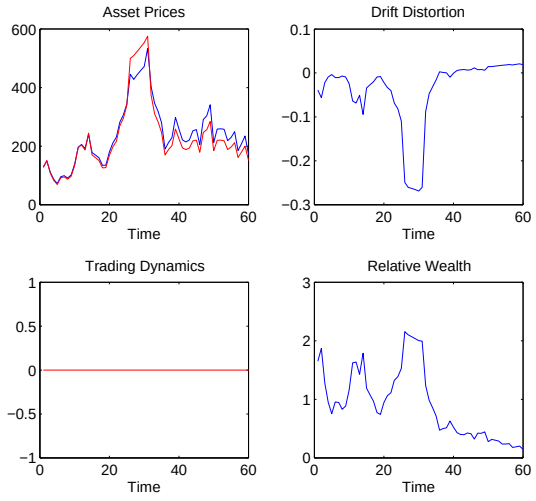

FIGURE 3. Simulation: $\mu = .01$, $\sigma^2 = .03$, $\varepsilon = .001$

Figure 3 reveals the interesting fact that 'bubbles' can emerge without any trade taking place. It is merely the *option* to trade that creates the bubble. Once again, the average detection error probability here is .47, implying that the agent's average doubts could not be statistically rejected by the historical data.

Finally, Figure 4 reports the results of a more atypical simulation. It is atypical for a couple of reasons. First, for these parameter values, we typically only see one or two trades. Three trades are rare. Second, the final trade is triggered by a huge drift distortion, in excess of 100%! Although this seems implausible, it is noteworthy that the average detection error probability using the above Chernoff bound only falls to 0.28. This suggests that we need to compute a tighter bound.

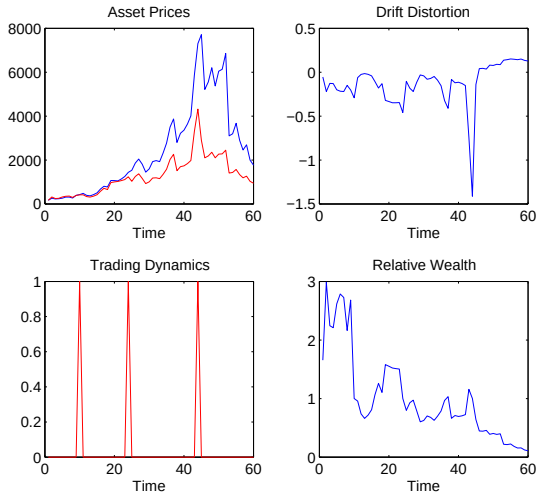

FIGURE 4. Simulation: $\mu = .01$, $\sigma^2 = .03$, $\varepsilon = .001$

6. CONCLUSION

TO BE ADDED

References

Anderson, E. W. (2005): "The Dynamics of Risk-Sensitive Allocations," *Journal of Economic Theory*, 125, 93–150.

Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003): "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection," *Journal of the European Economic Association*, 1, 68–123.

——— (2012): "Small Noise Methods for Risk-Sensitive/Robust Economies," *Journal of Economic Dynamics and Control*, 36, 468–500.

Barillas, F., L. P. Hansen, and T. J. Sargent (2009): "Doubts or Variability?," *Journal of Economic Theory*, 144, 2388–418.

Bhandari, A. (2013): "Doubts, Asymmetries, and Insurance," New York University.

Blume, L. E., T. Cogley, D. A. Easley, T. J. Sargent, and V. Tsyrennikov (2014): "The Case for Incomplete Markets," Cornell University.

Borovicka, J. (2013): "Survival and Long-Run Dynamics with Heterogeneous Beliefs Under Recursive Preferences," New York University.

Brunnermeier, M. K., A. Simsek, and W. Xiong (2014): "A Welfare Criterion for Models with Distorted Beliefs," *Quarterly Journal of Economics*, 129, 1753–97.

Duffie, D., and L. G. Epstein (1992): "Stochastic Differential Utility," *Econometrica*, 60, 353–94.

Froot, K. A., and M. Obstfeld (1991): "Intrinsic Bubbles: The Case of Stock Prices," *American Economic Review*, 81, 1189–1214.

Garber, P. (2000): *Famous First Bubbles: The Fundamentals of Early Manias*. MIT Press.

Gilboa, I., L. Samuelson, and D. Schmeidler (2014): "No-Betting Pareto Dominance," *Econometrica*, 82, 1405–42.

Hansen, L. P., and T. J. Sargent (2008): *Robustness*. Princeton University Press.

Hansen, L. P., T. J. Sargent, G. A. Turmuhambetova, and N. Williams (2006): "Robust Control, Min-Max Expected Utility, and Model Misspecification," *Journal of Economic Theory*, 128, 45–90.

Harrison, J. M., and D. M. Kreps (1978): "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics*, 92, 323–36.

Hong, H. J., and J. A. Scheinkman (2008): "Advisors and Asset Prices: A Model of the Origin of Bubbles," *Journal of Financial Economics*, 89, 268–87.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006): "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences," *Econometrica*, 74, 1447–98.

Maenhout, P. J. (2004): "Robust Portfolio Rules and Asset Pricing," *Review of Financial Studies*, 17, 951–83.

Morris, S. (1996): "Speculative Investor Behavior and Learning," *Quarterly Journal of Economics*, 111, 1111–33.

Santos, M., and M. Woodford (1997): "Rational Asset Pricing Bubbles," *Econometrica*, 65, 19–57.

Scheinkman, J. A. (2014): *Speculation, Trading, and Bubbles*. Columbia University Press.

Scheinkman, J. A., and W. Xiong (2003): "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 111, 1183–1219.

Shiller, R. J. (1989): *Market Volatility*. MIT Press.

Strzalecki, T. (2011): "Axiomatic Foundations of Multiplier Preferences," *Econometrica*, 79, 47–73.